



DIRECTORATE FOR SCIENCE, TECHNOLOGY AND INDUSTRY
COMMITTEE FOR INFORMATION, COMPUTER AND COMMUNICATIONS POLICY

DSTI/ICCP/CISP(2007)9
Unclassified

Working Party on Communication Infrastructures and Services Policy

GOOD PRACTICES IN INTERNET EXCHANGE POINT DOCUMENTATION AND MEASUREMENT

London, 24 - 25 May 2007

This paper was prepared by Bill Woodcock, Research Director at Packet Clearing House, and is not reserved for official use. It is for discussion under Item 11.

The paper sets out an initiative and methodology, produced within the Internet community, to improve measurement and documentation at IXPs, making more and better information available to all stakeholders. It is made available to delegates for information and comment. Delegates are invited to consider whether OECD governments could express support for or endorse this initiative in the broader context of encouraging the Internet community to improve transparency and make critical information on size and growth publicly available.

Contact: Bill Woodcock; e-mail is woody@pch.net

English - Or. English

TABLE OF CONTENTS

GOOD PRACTICES IN INTERNET EXCHANGE POINT DOCUMENTATION AND MEASUREMENT3

 Goals4

 Publication of Traffic Statistics.....5

 Publication of Participant Contacts.....5

 Maintenance of the IN-ADDR5

 This Document Describes Minimums5

 Summary and Next Steps6

FIGURE 16

GOOD PRACTICES IN INTERNET EXCHANGE POINT DOCUMENTATION AND MEASUREMENT

1. The Internet is a network of networks which has grown apace since the original NSFNET backbone was transitioned to a commercial environment. In March 2007 there were nearly 25 000 Autonomous Systems (i.e. networks with their own independent routing policy) present in the Internet Routing Table.¹ Unlike past communication networks, such as telegraph or circuit switched telephone networks, there are little data providing a holistic overview of the size and growth of Internet traffic around the world.
2. Many Autonomous Systems do collect some information on their own network and how it exchanges traffic with its immediate neighbours. Some go further and analyze their traffic exchange more generally, with the rest of the Internet. For the most part these data are regarded as proprietary and, therefore, not available to inform the many uses to which such information has been put in the past. Examples might include researchers developing technologies to support the carriage and exchange of traffic, capital markets or operators making decisions on demand for services and potential returns from investment in existing or new markets, aid or development agencies funding infrastructure in developing countries as well as the many uses such data could be put to in respect to informing policy making. This could range from an assessment of the affect of major outages and restoration, following natural disasters or malicious attacks, through to providing a more general picture of the size and growth of the Internet in any particular country following a change in regulation such as market liberalisation.
3. For governments such data would also enable progress in the implementation of the Tunis Agenda which recommended the establishment of Internet Exchange Points (IXPs) in the 92 countries which do not yet have them (Figure 1, Table 1).²
4. IXPs are of particular interest, as they are the principal potential sources of information on the development and growth of Internet traffic; critical information which has been unavailable since the deprecation of the centralized NSFnet backbone. IXPs are places where autonomous systems (e.g. Internet Service Providers or ISPs) exchange traffic between their networks. Since 1994 Packet Clearing House (PCH) has maintained a global directory of IXPs at: www.pch.net/ixpdir. In April 2007 this database listed the nearly 300 large and small IXPs around the world. The European Internet Exchange Association compiles a similar directory for its region with 111 IXPs listed in April 2007. This database can be viewed at: www.euro-ix.net/resources/list-eur.php.
5. Many IXPs already publish data on the scale and extent of their operations. These data have been collected on an ad hoc basis by both PCH and Euro-IX and made available to the Internet community through reports and presentations. Both organisations have, however, been increasingly approached by researchers, NGOs, and governmental and inter-governmental bodies, asking for a more complete and more accurate set of numbers which they can use to characterize the size and rate of growth of the Internet. In addition, at the NSF-OECD series of meetings on the “Future of the Internet”, the lack of data providing a holistic overview of Internet traffic was identified as a major deficiency in terms of better understanding inter-networking and the challenges and opportunities which are available in terms of future development.³

6. The lack of a standard approach to IXP documentation and measurement, harmonised across IXPs around the world, is an obstacle to meeting such requests. Accordingly, the goal of this document is to set out “good practices” in the documentation of Internet exchange points in a way which will allow their global adoption and implementation. It should be noted that, while external demand for statistics are growing, it is expected that the primary beneficiaries will be members of the IXPs themselves, as the value of an IXP to its members is directly dependent upon the number and size of ISPs participating, and clear documentation of the present value of an exchange is the best way to attract new participants, further increasing the amount of Internet capacity which it is able to create, in a virtuous cycle.

7. This document takes as a starting point ‘good practices’ in IXP documentation and measurement that have been developed around the world but particularly in Europe. To date, European IXPs have generally been most consistent and diligent in documenting their operations. This document takes current practices and distils common approaches into a harmonised methodology for data reporting. It notes a number of IXPs which are implementing the proposed methodology to provide experience and feedback. This methodology has been arrived at collectively through discussion and agreement in the first few months of 2007, among members of Euro-IX (the European Internet Exchange Point operators’ industry association), the AP IX SIG (the Special Interest Group of IXP operators in the Asia-Pacific region), NAPLA (the IXP operators group of Latin America), and other experts running IXPs around the world. It represents a global consensus on good practices within the industry.

8. An initial list of IXPs which are trialling the specific methodology documented herein are:

- Amsterdam Internet Exchange, The Netherlands
- Seattle Internet Exchange, United States
- Nepal Internet Exchange, Nepal
- Toronto Internet Exchange, Canada
- West Australia Internet Exchange, Australia
- Wellington Internet Exchange, New Zealand
- Internet Exchange of Puerto Rico

Goals

9. Specifically the document proposed documentation and measurement in three primary areas:

- Collection and publication of traffic statistics: The amount of Internet capacity being created at each IXP, measured at five-minute intervals.
- Publication of membership contacts: The list of Internet Service Providers that are collaborating to create capacity at each IXP, including contact information.
- Maintenance of the IN-ADDR zones: The machine-readable information about the ISPs which allows diagnostic tools like Traceroute and Dig to accurately identify ISPs’ networks.

10. It was deemed critical that the documented practices consist exclusively of techniques which were already in widespread use, and widely demonstrated to be easily implemented. This document is intended to specify readily achievable goals, with an eye to universal or essentially-universal adoption. The document harmonizes a broad and widely divergent set of practices by boiling them down to a simple and uniform subset. The results are intended to be both human-readable and machine-readable, with a particular emphasis on making results accessible and intelligible despite any lack of a shared language between the researcher seeking data and the IXP operator publishing it.

Publication of Traffic Statistics

11. On the IXP web site, in addition to any other localized or translated Uniform Resource Locator (URLs), a page shall be located at /statistics/ relative to the root of the web site, containing any statistical information which the IXP chooses to publish, in the language of its choice. This shall include at a minimum, a graphical histogram of the past 24 hours' aggregate traffic, counting each byte which has flowed across the IXP switch fabric exactly once, with care not to double-count at inter-switch connections. The time-base of the graph may be in the local time-zone of the IXP. The graph image itself shall also be reachable at /statistics/graph/. The same data shall also be provided in tabular form at /statistics/table/, of at least 288 lines, where each consists of a Unix-style timestamp five-minute aligned Coordinated Universal Time (UTC), a tab, and the number of bits passed in the preceding five-minute period, divided by 300 and rounded to the nearest integer (in other words, five-minute average bits per second), followed by a carriage return, and the next entry. In the event that data is simply unavailable, no value should follow the tab. Under no circumstances should a lack of data be represented by a zero or an interpolated value. The earliest entry shall be at the beginning (top) of the file, and the most recent entry shall be at the end (bottom) of the file.

Publication of Participant Contacts

12. On the IXP web site, in addition to any other localized or translated URLs, a page shall be located at /participants/ relative to the root of the web site, containing any information about the IXP's membership or participants which the IXP chooses to publish, in the language of its choice. This shall include, at a minimum, the organization name and contact method, such as a peering information URL, for each participant. Participant data shall also be provided in tabular form at /participants/table/, exactly one line per assigned IP address of each subnet used for public interconnection across the IXP, where each line consists of the assigned IP address, a tab, the organization name of the participant to which it has been assigned, a tab, the principal Autonomous System Number which the participant reports to be in use on that interface, a tab, and the preferred contact information for the responsible party for that participant organisation, followed by a carriage return, and the next entry. Unassigned IP addresses shall not appear in the file. The file shall be sorted by IP address, with the lowest address at the beginning (top) of the file, and the highest address at the end (bottom) of the file.

Maintenance of the IN-ADDR

13. In the publicly-visible IN-ADDR Domain Name System (DNS), one PTR record⁴ shall be maintained for each assigned IP address in each subnet used for public interconnection across the IXP, consisting of the fully-qualified domain name received from the participant to which that IP address is assigned, and each participant shall be encouraged to maintain a matching A record in their forward domain. Unassigned IP addresses shall have no associated PTR record in the publicly-visible DNS. Address ranges for which the IXP does not control authoritative IN-ADDR delegation shall not be used for the exchange of Internet traffic.

This Document Describes Minimums


14. It is important to understand that this document describes a least common denominator practice, intended to facilitate easy programmatic interaction, and is in no way intended to discourage the implementation of additional features or language localization for the convenience of the IXP's first constituency, its participants. Nor does the document suggest precluding other approaches in respect to Internet measurements by communities other than IXPs through, for example, co-operation between researchers and network operators. What the suggested approach here aims to achieve is to harmonise data collection and reporting by IXPs as a resource for IXP members as well as other stakeholders.

Summary and Next Steps

15. This document out an initiative and methodology, produced within the Internet community, to improve measurement and documentation at IXPs, making more and better information available to all stakeholders.

16. PCH would welcome comments by OECD member countries and the distribution of this proposal to IXPs in your countries. Delegates are invited to consider whether OECD governments could express support for or endorse this initiative in the broader context of encouraging the Internet community to improve transparency and make critical information on size and growth publicly available.

FIGURE 1



Density of IXPs

TABLE 1

Countries with IXPs: 78	Countries without IXPs: 92
AR	1 AD
AT	1 AE
AU	9 AF
BD	1 AL
BE	2 AM
BR	10 AO
BW	1 AZ
CA	4 BA
CD	1 BF
CH	3 BG
CN	1 BI
CO	1 BJ
CU	1 BN
CY	1 BO
CZ	1 BT
DE	12 BY
DK	1 BZ
EC	2 CF
EE	2 CG
EG	1 CI
ES	6 CL
FI	2 CM
FR	7 CR
GB	12 DJ
GH	1 DO
GR	1 DZ
HK	1 EH
HR	1 ER
HU	1 ET
ID	5 GA
IE	1 GE
IL	1 GF
IN	4 GL
IS	1 GM
IT	4 GN
JP	17 GQ
KE	1 GT
KH	1 GW
KR	5 GY
LK	1 HN
LU	1 HT
LV	1 IQ
MN	1 IR
MT	1 JM
MU	1 JO
MY	1 KG
MZ	1 KP
NG	1 KW
NI	1 KZ
NL	4 LA
NO	1 LB
NP	1 LI
NZ	4 LR
PE	2 LS
PH	4 LT
PK	1 LY
PL	1 MA
PR	1 MD
PT	1 MG
PY	1 MK
RO	2 ML
RU	7 MM
SA	1 MR
SE	9 MW
SG	2 MX
SI	1 NA
SK	2 NE
SZ	1 OM
TH	1 PA

TABLE 1 (cont'd.)

Countries with IXPs: 78	Countries without IXPs: 92
TR	1 PG
TW	4 PS
TZ	1 QA
UA	2 RS
UG	1 RW
US	77 SD
VN	2 SL
ZA	3 SN
ZW	1 SO
	SR
	SV
	SY
	TD
	TG
	TJ
	TM
	TN
	TL
	UY
	UZ
	VE
	YE
	ZM

Source: PCH (Accessed 25 April 2007) <http://www.pch.net/ixpdir/summary>

¹ <http://www.apnic.net/mailling-lists/bgp-stats/>

² paragraph 50 sub points b and c. “Tunis Agenda for the Information Society”, WSIS-05/TUNIS/DOC/6(Rev. 1)-E , 18 November 2005 <http://www.itu.int/wsis/docs2/tunis/off/6rev1.html>

³ SF/OECD Workshop “Social and Economic Factors Shaping the Future of the Internet”, 31 January 2007 http://www.oecd.org/document/59/0,2340,en_2649_34223_37921851_1_1_1_1,00.html and OECD ICCP Workshop: “The Future of the Internet”, Paris, 8 March 2006 http://www.oecd.org/document/5/0,2340,en_2649_34223_36169989_1_1_1_1,00.html

⁴ PTR record or pointer record maps an IPv4 address to the canonical name for that host. Setting up a PTR record for a hostname in the in-addr.arpa domain that corresponds to an IP address implements reverse DNS lookup for that address. http://en.wikipedia.org/wiki/Domain_Name_System